arXiv:1703.03453v1 [cs.AI] 9 Mar 2017

Using Options for Long-Horizon Off-Policy Evaluation

Zhaohan Daniel Guo[1], Philip S. Thomas[1], Emma Brunskill[2]

Abstract

Evaluating a policy by deploying it in the real world can be risky and costly. *Off-policy evaluation* (OPE) algorithms use historical data collected from running a previous policy to evaluate a new policy, which provides a means for evaluating a policy without requiring it to ever be deployed. *Importance sampling* is a popular OPE method because it is robust to partial observability and works with continuous states and actions. However, we show that the amount of historical data required by importance sampling can scale exponentially with the *horizon* of the problem: the number of sequential decisions that are made. We propose using policies over temporally extended actions, called *options*, to address this long-horizon problem. We show theoretically and experimentally that combining importance sampling with options-based policies can significantly improve performance for long-horizon problems.

1. Introduction

One important problem for many high-stakes sequential decision making under uncertainty domains, including robotics, healthcare, education, and dialogue systems, is estimating the performance of a new policy without requiring it to be deployed. To address this, off-policy evaluation (OPE) algorithms use historical data collected from executing a behavior policy, to predict the performance of the new policy (the evaluation policy). Importance sampling (IS) is one powerful approach that can be used to evaluate the potential performance of a new policy for a sequential decision making problem under uncertainty (Precup, 2000). In contrast to model based approaches to OPE (e.g. Hallak et al. (2015a)), importance sampling provides an unbiased estimate of the evaluation policy, regardless of the accuracy of the input domain model representation. In particular, this makes importance sampling robust to partial observability, which is often prevalent in specifications of real-world domains. Unfortunately, importance sampling estimates of the evaluation policy can be very inaccurate when the horizon of the problem is long: the variance of IS estimators can grow exponentially with the number of sequential decisions made in an episode. This is a serious limitation for many applications which involve decisions made over tens or hundreds of steps (consider dialogue systems or educational software) like intelligent tutoring systems, which must make dozens of decisions about how to sequence the content shown to a student, and dialogue systems where a conversation might require dozens of responses).

Due to the importance of OPE, there have been many recent efforts to improve the accuracy of importance sampling. For example, Dudík et al. (2011) and Jiang & Li (2016) proposed doubly robust importance sampling estimators that can greatly reduce the variance of predictions when an approximate model of the environment is available. Thomas & Brunskill (2016) proposed an estimator that further integrates importance sampling and model-based approaches, and which can greatly reduce mean squared error. These approaches trade between the bias and variance of model-based and importance sampling approaches, and result in consistent estimators. Unfortunately, in long horizon settings, these approaches will either create estimates that suffer from high variance or may end up exclusively relying on the provided approximate model, potentially inducing high bias. Other recent efforts that achieve improved OPEs estimate a *value function* using off-policy data rather than just the performance of a policy (Munos et al., 2016; Hallak et al., 2015b); however, such models again suffer from bias if the input state description is not Markov (such as if the domain description induces partial observability).

To provide better off policy estimates in long horizon domains, we propose leveraging temporal abstraction. In particular, we analyze using options-based policies (policies with temporally extended actions) (Sutton et al., 1999) instead of policies over primitive actions. We prove that the we can obtain an exponential reduction in the variance of the resulting estimates, and in some cases, cause the variance to be independent of the horizon. We also demonstrate this benefit with simple simulations. Crucially, our results can be equivalently viewed as showing that using options can drastically reduce the amount of historical data required to obtain an accurate estimate of a new evaluation policy's performance.

[1]Carnegie Mellon University, Pittsburgh, PA
[2]Stanford University, Stanford, CA

2. Background

Here we provide a brief background into reinforcement learning, options and OPE. We consider sequential stochastic decision processes in which at each step an agent can select an action, and then receives a reward, and a new observation: $s_t^{(i)} \in S$ is a discrete state, $a_t^{(i)} \in A$ is a discrete action, and $r_t^{(i)}$ is a real-valued reward in $[0, R_{\max}]$. The dynamics and reward are unknown and are denoted by the transition probability $T(s_{t+1}^{(i)}|s_1^{(i)}, a_1^{(i)}, \ldots, s_t^{(i)}, a_t^{(i)})$ and reward density $R(r_t^{(i)}|s_t^{(i)}, a_t^{(i)})$. A primitive policy maps histories to action probabilities i.e. $\pi(a_t|s_1, a_1, \ldots, s_t)$ is the probability of executing action a_t after history $s_1, a_1, \ldots, s_t$. The return of a trajectory of H steps is simply the sum of the rewards $J(\tau) = \sum_{t=1}^{H} r_t$. The value of a policy π is the expected return when running that policy $V_\pi = \mathbb{E}_\pi(J(\tau))$.

Temporal abstraction can help speed planning and online learning (Sutton et al., 1999; Mann & Mannor, 2013; Mankowitz et al., 2014; Brunskill & Li, 2014). One popular form of temporal abstraction are sub-policies, in particular options (Sutton et al., 1999). An option o consists of a primitive policy π (a policy over regular actions), a termination condition $\beta : \Omega \to [0, 1]$ where $\beta(\tau)$ is the probability of stopping the option given the current trajectory $\tau \sim \Omega$ from when this option began, and an input set $I \subset S$ where $s \in I$ denotes the states where option o is allowed to start. Primitive actions can be considered as a special case of options that always terminate after a single step. An options-based policy $\mu(o_t|s_1, a_1, \ldots, s_t)$ denotes the probability of picking option o_t given history $(s_1, a_1, \ldots, s_t)$ when the previous option has terminated. A high-level trajectory of length k is denoted by $T = (s_1, o_1, v_1, s_2, o_2, v_2, \ldots, s_k, o_k, v_k)$ where v_t is the sum of the rewards accumulated when executing option o_t.

In this paper we will consider batch, offline, off-policy evaluation of policies for sequential decision making domains using both primitive action policies and options-based policies. We will now introduce the general OPE problem using primitive policies: in a later section we will combine this with options-based policies.

In OPE we assume access to historical data D generated by a stochastic decision process and a behavior policy π_b. D consists of n trajectories $\{\tau^{(i)}\}_{i=1}^n$. A trajectory has length H (i.e. the horizon) and is $\tau^{(i)} = (s_1^{(i)}, a_1^{(i)}, r_1^{(i)}, s_2^{(i)}, a_2^{(i)}, r_2^{(i)}, \ldots, s_H^{(i)}, a_H^{(i)}, r_H^{(i)})$. In off-policy evaluation, the goal is to use the data D to estimate the value of an evaluation policy π_e: V_{π_e}. As D was generated from running the behavior policy π_b, we cannot simply use the Monte Carlo estimate. An alternative is to use importance sampling to reweight the data in D to give greater weight to samples that are likely under π_e and lesser weight to unlikely ones. We consider per-decision importance sampling (PDIS) (Precup, 2000), which gives the following estimate of the value of π_e:

$$PDIS(D) = \frac{1}{n} \sum_{i=1}^{n} \left(\sum_{t=1}^{H} \rho_t^{(i)} r_t^{(i)} \right) \tag{1}$$

$$\rho_t^{(i)} = \prod_{u=1}^{t} \frac{\pi_e(a_u^{(i)}|s_u^{(i)})}{\pi_b(a_u^{(i)}|s_u^{(i)})} \tag{2}$$

where $\rho_t^{(i)}$ is the weight given to the rewards to correct for the difference in distribution. This estimator gives us an unbiased estimate of the value of π_e:

$$\mathbb{E}_{\pi_e}(J(\tau)) = \mathbb{E}_{\pi_b}(PDIS(\tau)) \tag{3}$$

For simplicity from here onwards we will assume that primitive and options-based policies are a function only of the current state, but our results apply also when the they are a function of the history. Note that importance sampling does not assume that the states in the trajectory are Markovian, and is thus robust to error in the state representation, and in general, robust to partial observability as well.

3. Importance Sampling and Long Horizons

We now show how the amount of data required for importance sampling to obtain a good off-policy estimate can scale exponentially with the problem horizon. Notice that in the standard importance sampling estimator, the weight is the product of the ratio of action probabilities. We now prove that this can cause the variance of the policy estimate to be exponential in H.

Theorem 1. *The mean squared error of the PDIS estimator can be* $\Omega(2^H)$.

Proof. Because PDIS is an unbiased estimator of an evaluation policy's performance, its MSE is equal to its variance. To prove the theorem statement, we provide an existence proof by constructing a sample MDP where, given a particular behavior policy, there is an evaluation policy whose estimate under PDIS will have a variance that scales exponentially with the horizon H.

Consider a discrete state and action Markov decision process. The horizon is H and the MDP has $2H + 1$ states and 2 actions. The states form two chains: the top chain has length $H + 1$ and the bottom chain has length H. Label the states of the top chain as $x_1, \ldots, x_{H+1}$, and the states of the bottom chain be $y_1, \ldots, y_H$. The start state is x_1. An episode halts after H steps. The two actions are a_1, a_2. Taking action a_1 in the top chain deterministically transitions to the next state in the top chain i.e. from x_i to x_{i+1}. Taking action a_2 in the top chain determinitically

transitions to the corresponding state in the bottom chain i.e. from x_i to y_i. The reward is zero everywhere except a reward of 1 is received for executing action a_1 at state x_H. The optimal policy is to always take action a_1.

Let the behavior policy π_b be uniformly random i.e. there is always a probability of $1/2$ of picking either action. The evaluation policy is the optimal policy, $\pi_e(s) = a_1$ for all states.

Since the only nonzero reward is the single reward of 1 at x_H for action a_1, and it is only possible to reach that state by taking action a_1 for H steps, PDIS reduces to a sum only over trajectories consisting solely of H steps of action a_1, whose weights are $\rho = \prod_{u=1}^{H} \frac{\pi_e(a_1|s_u^{(i)})}{\pi_b(a_1|s_u^{(i)})} = 2^H$. The PDIS estimate of the evaluation policy is a scaled Binomial distribution where with probability $p = \frac{1}{2^H}$ a trajectory's weighted return is 2^H and zero otherwise. Thus the variance of the PDIS estimate of π_e is $\frac{1}{n}(2^H - 1)$ for n trajectories, which is $\Omega(2^H)$. □

Equivalently, this means that achieving a desired mean squared error of ϵ can require a number of trajectories that scales exponentially with the horizon.

A natural question is whether this issue also arises in a weighted importance sampling(Precup et al., 2001), a popular (biased) approach to OPE that has lower variance. We show below that the long horizon problem still persists.

Theorem 2. *It takes $\Omega(2^H)$ trajectories to linearly reduce the MSE of weighted importance sampling (WIS) in the worst case.*

Proof. We prove the above statement by constructing a MDP and selecting a behavior and target policy which will result in the stated MSE dependence on the horizon. We consider the same MDP, π_b, and π_e as used in the proof of Theorem 1. For this particular MDP, the only weight that matters is the weight associated with the single final reward of the correct trajectory, so per-decision importance sampling and ordinary importance sampling are equivalent.

If the optimal trajectory does not appear in the historical data, then WIS is undefined. This is because the weight of any nonoptimal trajectory is 0, so dividing by the sum of the weights is undefined. However if we change π_e from deterministically picking a_1 to picking a_1 with arbitrarily high probability, then the weights of nonoptimal trajectories will be arbitrarily close to zero, resulting in a WIS estimate of 0. Thus we define WIS to estimate a value of 0 when the optimal trajectory does not appear in the data. Any optimal trajectory that appears in your data will have a weight of 2^H. Then because the weights of nonoptimal trajectories are 0, the WIS estimate will be exactly 1. Thus as soon as WIS sees at least one correct trajectory it will have the perfect estimate, otherwise the estimate will be 0. The WIS estimate is a Bernoulli distribution where the probability of 1 is the probability of at least one optimal trajectory appearing in the data.

Since the WIS estimate is Berounlli, its variance is bounded by a constant. Furthermore the variance is small. Thus we take a closer look at the bias, since MSE is the sum of the variance and bias squared. First we compute the probability the WIS returns 1. This is the probability of at least one optimal trajectory appearing, which is equivalent to one minus the probability of no optimal trajectory appearing: $1 - \left(1 - \frac{1}{2^H}\right)^n$. Thus the expected value of the WIS estimate is $1 - \left(1 - \frac{1}{2^H}\right)^n$. Then the bias is $\left(1 - \frac{1}{2^H}\right)^n$. Let the bias be B. We will compute how much data is needed to compensate for the increase in the bias when H increases. Rearranging and solving for n (using a taylor approximation) we get $n = \frac{\log B}{\log\left(1 - \frac{1}{2^H}\right)} \approx \frac{\log B}{-\frac{1}{2^H}} \approx \Omega(2^H)$. Thus we need an exponential number of trajectories to compensate for the increase in bias when the horizon H is increased. Since MSE consists partly of biased squared, we would need even more data to compensate for a squared increase in bias, but for simplicity we still use an exponential bound.

□

4. Combining Options and Importance Sampling

We will show that we can take advantage of options to help mitigate the issue of a long horizon. If the behavior and evaluation policy are both options-based policies, then the PDIS estimator can be exponentially more data efficient compared to using primitive behavior and evaluation policies.

Due to the structure in options-based policies, we can decompose the difference between the behavior policy and the evaluation policy in a natural way. Let μ_b be the options-based behavior policy and μ_e be the options-based evaluation policy. First, we examine the probabilities over the options. The probabilities $\mu_b(o_t|s_t)$ and $\mu_e(o_t|s_t)$ can differ and contribute a ratio of probabilities as an importance sampling weight. Second, the underlying policy π for an option o_t present in both μ_b and μ_e may differ, and this also contributes to the importance sampling weights. Finally, additional or missing options can be expressed by setting the probabilities over missing options to be zero for either μ_b or μ_e. Using this decomposition, we can easily apply PDIS to options-based policies.

Theorem 3. *Let $\mathcal{O}$ be the set of options that have the same underlying policies between μ_b and μ_e. Let $\overline{\mathcal{O}}$ be the set of options that have changed underlying policies. Let $k^{(i)}$ be the length of the i-th high level trajectory from data set*

D. Let $j_t^{(i)}$ be the length option $o_t^{(i)}$. The PDIS estimator applied to D is

$$PDIS(D) = \frac{1}{n}\sum_{i=1}^{n}\left(\sum_{t=1}^{k^{(i)}} w_t^{(i)} y_t^{(i)}\right) \tag{4}$$

$$w_t^{(i)} = \prod_{u=1}^{t}\frac{\mu_e(o_u^{(i)}|s_u^{(i)})}{\mu_b(o_u^{(i)}|s_u^{(i)})} \tag{5}$$

$$y_t^{(i)} = \begin{cases} v_t^{(i)} \text{ if } o_t^{(i)} \in \mathcal{O} \\ \sum_{b=1}^{j_t^{(i)}} \rho_{t,b}^{(i)} r_{t,b}^{(i)} \text{ if } o_t^{(i)} \in \overline{\mathcal{O}} \end{cases} \tag{6}$$

$$\rho_{t,b}^{(i)} = \prod_{c=1}^{j_t^{(i)}}\frac{\pi_e(a_{t,c}^{(i)}|s_{t,c}^{(i)}, o_t^{(i)})}{\pi_b(a_{t,c}^{(i)}|s_{t,c}^{(i)}, o_t^{(i)})} \tag{7}$$

where $r_{t,b}^{(i)}$ is the b-th reward in the subtrajectory of option $o_t^{(i)}$ and similarly for s and a.

Proof. This is a straightforward application of PDIS to the options-based policies using the decomposition mentioned. □

Theorem 3 expresses the weights as made up of two parts: one part comes from the probabilities over options which is expressed as $w_t^{(i)}$, and another part comes from the underlying primitive policies of options that have changed with $\rho_{t,b}^{(i)}$. We can immediately make some interesting observations below.

Corollary 1. *If no underlying policies for options are changed between μ_b and μ_e, and all options have length at least K steps, then the worst case variance of PDIS is exponentially reduced from $\Omega(2^H)$ to $\Omega(2^{(H/K)})$*

Corollary 1 follows from Theorem 3. Since no underlying policies are changed, then the only importance sampling weights left are $w_t^{(i)}$. Thus we can focus our attention only on the high-level trajectory which has length at most H/K. Effectively the horizon has shrunk from H to H/K which results in the exponential reduction of the worst case variance of PDIS. This result suggests that if we already have a set of options that we are satisfied with and do not wish to change, and we only care about deciding which option to pick, then we get the benefit of reduced variance from a small horizon.

Corollary 2. *If the probabilities over options are the same between μ_b and μ_e, and a subset of options $\overline{\mathcal{O}}$ have changed their underlying policies, then the worst case variance of PDIS is reduced from $\Omega(2^H)$ to $\Omega(2^K)$ where K is an upper bound on the sum of the lengths of the options.*

Corollary 2 follows from Theorem 3. The options whose underlying policies are the same between behavior and evaluation can effectively be ignored, and cut out of the trajectories in the data. This leaves only options whose underlying policies have changed, shrinking down the horizon from H to the length of the leftover options. For example, if only a single option of length 3 is changed, and the option appears once in a trajectory, then the horizon can be reduced to just 3. This result can be very powerful, as the reduced variance becomes independent of the horizon H.

The two scenarios above are not the only cases. With options-based policies, any mixture of changing probabilities over options and changing underlying policies can work. Theorem 3 can be used as a reference to see how many terms remain in the products of the weights to get a sense of what the effective horizon could be.

5. Going Further with Options

Often options are used to achieve a specific sub-task in a domain. For example in a robot navigation task, there may be an option to navigate to a special fixed location. However one may realize that there is a faster way to navigate to that location, so one may change that option and try to evaluate the new policy to see whether it is actually better. In this case the old and new option are both always able to reach the special location; the only difference is that the new option could get there faster. In such a case we can take advantage of this property to further reduce the variance of PDIS. We now formally define this property.

Definition 1. *Given behavior policy μ_b and evaluation policy μ_e, an option o is called **stationary**, if no matter where option o starts to execute, the distribution of termination states (states on which o terminates) is the same for μ_b and μ_e. The underlying policy for option o is allowed to be different between μ_b and μ_e; only the termination state distribution is important.*

A stationary option may not always arise due to solving a sub-task. It can also be the case that a stationary option is used as a way to perform a soft reset. For example, a robotic manipulation task may want to reset arm and hand joints to a default configuration in order to minimize sensor/motor error, before trying to grasp a new object.

Stationary options allows us to point to a step in a trajectory where we know the state distribution is fixed. Because the state distribution is fixed, we can actually partition the trajectory into two. The beginning of the second partition would then have a fixed state distribution. Then we can independently apply PDIS to each partition, and then sum up the estimates; we can perform the sum precisely because the starting state distribution of the second partition is fixed. This is extremely powerful as it allows us to partition hori-

zon of the problem.

We can extend this to when a task may have multiple points in a trajectory with a fixed stated distribution. This may come up when a task has multiple sub-tasks or multiple soft resets. Then we can partition the trajectories at each point, resulting in many small sub-trajectories, and apply PDIS to each part before summing the results. This can lead to an even greater reduction in the horizon.

Furthermore, we can make one more extension where we may have a choice of different stationary options that all have the same terminating state distribution. Perhaps in a task we have multiple ways to solving a sub-task or multiple ways to perform a soft reset.

Theorem 4 below summarizes the benefit from stationary options.

Theorem 4. *Let μ_b be an options-based behavior policy. Let μ_e be an options-based evaluation policy. Let $\mathcal{O}$ be the set of options that μ_b and μ_e use. The underlying policies of the options in μ_e may be arbitrarily different from μ_b.*

Let $\mathcal{O}_1 \subset \mathcal{O}$ be a subset of options that are all stationary with the same terminating state distribution. Similarly for $\mathcal{O}_2, \ldots, \mathcal{O}_l$. Let $\mathcal{O}' \subset \mathcal{O}$ be the subset of other options not part of any $\mathcal{O}_i$.

Let $(o_1, o_2, \ldots, o_l)$ denote a trajectory where o_i is chosen from $\mathcal{O}_i$).

Suppose all behavior trajectories τ follow the form of $(o_1, o_2, \ldots, o_l)$ with any number of options $o \in \mathcal{O}'$ interspersed.

Then, we can decompose the expected value as follows. Let τ_1 be the first part of a trajectory up until and including the first occurrence of o_1. Let τ_2 be the part of the trajectory after the first occurrence of o_1 up to and including the first occurrence of o_2. Similarly for τ_j. Then the expected value is

$$\mathbb{E}_{\mu_e}(J(\tau)) = \mathbb{E}_{\mu_b}(PDIS(\tau)) \tag{8}$$

$$= \sum_{j=1}^{l} \mathbb{E}_{\mu_b}(PDIS(\tau_j)) \tag{9}$$

Proof. Let t_j^* be the timestep when o_j terminates. Let $t_0^* = 1$ be the first timestep. Then

$$\mathbb{E}_{\mu_e}(J(\tau)) \tag{10}$$

$$= \mathbb{E}_{\mu_e}\left(\sum_{j=1}^{l} J(\tau_j)\right) \tag{11}$$

$$= \sum_{j=1}^{l} \mathbb{E}_{\mu_e}(J(\tau_j)) \tag{12}$$

$$= \sum_{j=1}^{l} \mathbb{E}_{\mu_e}\left(\mathbb{E}_{\mu_e}(J(\tau_j)|s_{t_{j-1}^*} = s)\right) \tag{13}$$

$$= \sum_{j=1}^{l} \sum_{s \in S} \Pr(s_{t_{j-1}^*} = s|\mu_e)\left(\mathbb{E}_{\mu_e}(J(\tau_j)|s_{t_{j-1}^*} = s)\right) \tag{14}$$

$$= \sum_{j=1}^{l} \sum_{s \in S} \Pr(s_{t_{j-1}^*} = s|\mu_e)\left(\mathbb{E}_{\mu_b}(PDIS(\tau_j)|s_{t_{j-1}^*} = s)\right) \tag{15}$$

$$= \sum_{j=1}^{l} \sum_{s \in S} \Pr(s_{t_{j-1}^*} = s|\mu_b)\left(\mathbb{E}_{\mu_b}(PDIS(\tau_j)|s_{t_{j-1}^*} = s)\right) \tag{16}$$

$$= \sum_{j=1}^{l} \mathbb{E}_{\mu_b}\left(\mathbb{E}_{\mu_b}(PDIS(\tau_j)|s_{t_{j-1}^*} = s)\right) \tag{17}$$

$$= \sum_{j=1}^{l} \mathbb{E}_{\mu_b}(PDIS(\tau_j)) \tag{18}$$

where eqn 13 follows from the law of total expectation, eqn 15 follows from using PDIS with a fixed initial state distribution $s_{t_{j-1}^*} = s$, eqn 16 follows because $s_{t_{j-1}^*}$ is the terminating state for option o_{j-1} whose terminating state distribution stayed the same between μ_b and μ_e, and eqn 18 follows from the law of total expectation. □

Corollary 3. *If the options in Theorem 4 have length at most K, then the worst case variance of PDIS is reduced from $\Omega(2^H)$ to $\Omega((H/K)2^K)$.*

Note that there are no conditions on how the probabilities over options may differ, nor on how the underlying policies of the non-stationary options may differ. This means regardless of these differences, the trajectories can be partitioned and PDIS can be independently applied. Furthermore, Theorem 3 can still be applied to each of the independent applications of PDIS. Combining Theorem 4 and Theorem 3 can lead to more ways of designing a desired evaluation policy that will result in a low variance PDIS estimate.

Note that the benefits of a fixed state distribution can also apply to primitive policies, as long as there is a consistent way to determine a point in a trajectory where the state distribution stays the same. However, this special case may not always be applicable, but by using options-based behavior and evaluation policies, this case can naturally arise in many situations with multiple sub-tasks or soft resets.

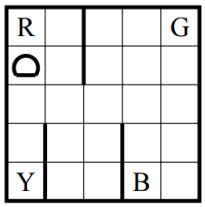

Figure 1. Taxi Domain (Dietterich, 2000)

6. Using a Primitive Behavior Policy

If the collected data is from a primitive policy, it is still possible in a restricted sense to apply the techniques for options-based policies. We can reinterpret the primitive behavior policy as a specific options-based policy. First, one is free to pick the number of options, but they all will have the primitive policy as the underlying policy. Second, the input set and termination conditions of the options can be freely chosen. Finally, the probabilities over the options can also be freely chosen as long as they are consistent with the input sets.

7. Experiments

Our experiments serve as a simple illustration of how viewing things through the options framework can significantly improve the accuracy of importance-sampling-based off policy estimation in long horizon domains. Since importance sampling is particularly useful when a good model of the domain is unknown and/or the domain involves partial observability, we introduced a partially observable variant of the popular Taxi domain (Dietterich, 2000) for our simulations.

7.1. Partially Observable Taxi

First we describe the standard fully observable Taxi domain(Dietterich, 2000). It is a 5x5 gridworld (Figure 7.1). There are 4 special locations: R,G,B,Y. A passenger starts randomly at one of the 4 locations, and its destination is randomly chosen from one of the 4 locations. The taxi starts randomly on any square. The taxi can move one step in any of the 4 cardinal directions N,S,E,W, as well as attempt to pickup or drop off the passenger. Each step has a reward of -1. An invalid pickup or dropoff has a -10 reward and a successful dropoff has a reward of 20.

In our NoisyTaxi, the location of the taxi and the location of the passenger is partially observable. If the row location of the taxi is c, the agent observes c with probability 0.85, $c+1$ with probability 0.075 and $c-1$ with probability 0.075 (if adding or subtracting 1 would cause the location to be outside the grid, the resulting location is constrained to still lie in the grid). The column location of the taxis is observed with the same noisy distribution. Before the taxi successfully picks up the passenger, the observation of the location of the passenger has a probability of 0.15 of switching randomly to one of the four designated locations. After the passenger is picked up, the passenger is observed to be in the taxi with 100% probability (e.g. no noise while in the taxi).

7.2. Experiment 1: Options-based Policies

We first illustrate how using options-based policies can yield significantly tighter estimates of new policies (that use the same set of options but may have a different policy over them) compared to using primitive policies.

We consider ϵ-greedy option policies, where with probability $1-\epsilon$ the policy samples the optimal option, and probability ϵ the policy samples a random option. Options in this case are n-step policies, where "optimal" options involve taking n-steps of the optimal (primitive action) policy, and "random" options involve taking n random primitive actions. Our behavior policies π_b will use $\epsilon = 0.3$ and our evaluation policies π_e use $\epsilon = 0.05$. We investigate how the accuracy of estimating π_e varies as a function both of the number of trajectories and the length of the options $n = 1, 2, 3$. Note $n = 1$ corresponds to having a primitive action policy.

Empirically, all behavior policies have essentially the same performance. Similarly all evaluation policies have essentially the same performance. We first collect data using the behavior policies, and then use PDIS to evaluate their respective evaluation policies.

Our results show that PDIS for the options-based evaluation policies are an order of magnitude better than PDIS for the primitive evaluation policy. Figure 7.2 shows the MSE (log scale) of the PDIS estimators for the evaluation policies. Indeed, Corollary 1 shows that the n-step options policies are effectively reducing the horizon by a factor of n over the primitive policy. As expected, the options-based policies that use 3-step options have the lowest MSE. This shows the substantial variance reduction benefits of using options-based policies for both the behavior and evaluation policy.

7.3. Experiment 2: Stationary Options

In Theorem 4 we showed that if many options end with the same distribution over terminal states between behavior and evaluation policies, then this could be leveraged to break apart long products of importance weights into additive sums of shorter products. This can yield substantial improvements in the resulting variance over the estimated evaluation policy value.

We now demonstrate this empirically. Interestingly we

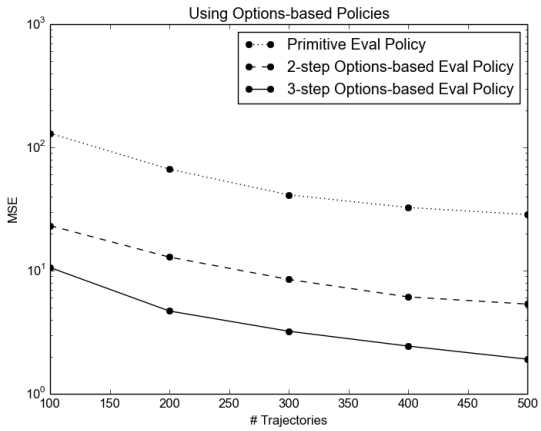

Figure 2. Comparing the MSE of PDIS between primitive and options-based behavior and evaluation policy pairs. Note the y-axis is a log scale.

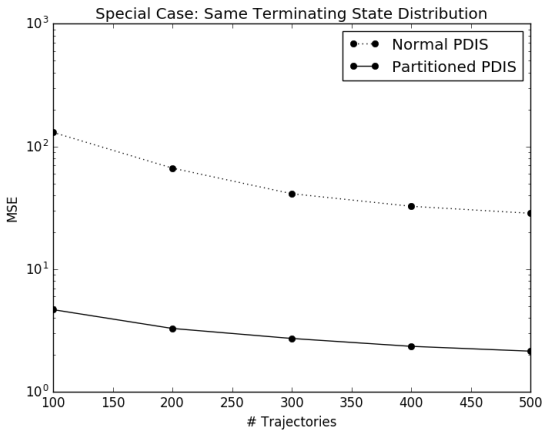

Figure 3. Comparing MSE of Normal PDIS and PDIS that uses Theorem 4

don't have to explicitly construct options for NoisyTaxi: rather we just have to identify points in the trajectory in which the behavior and evaluation policy would reach an identical distribution over states. This allows us to break the trajectory into subparts.

More specifically, in NoisyTaxi, we know that a primitive ϵ-greedy policy will eventually pick up the passenger (though it may take a very long time depending on ϵ). Since the starting location of the passenger is uniformly random, the location of the taxi immediately after picking up the passenger is also uniformly random. This implies that regardless of the ϵ value in an ϵ-greedy policy, that we can view executing that ϵ-greedy policy until the passenger is picked up as a new "PickUp-ϵ" option that always terminates in the same state distribution.

Given this argument, we can use Theorem 4 to decompose any NoisyTaxi trajectory into the part before the passenger is picked up, and the part after the passenger is picked up, estimate the expected reward for each, and then sum. As picking up the passenger is often the halfway point in a trajectory (depending on the locations of the passenger and the destination), we can perform importance sampling over two, approximately half length, trajectories.

More concretely, we consider two $n = 1$ options (e.g. primitive action) ϵ-greedy policies. Like in the prior subsection, the behavior policy has $\epsilon = 0.3$ and the evaluation policy has $\epsilon = 0.05$. We compare performing normal PDIS to estimate the value of the evaluation policy to estimating it using partitioned-PDIS using Theorem 4.

We gain an order of an order of magnitude reduction in MSE (labeled Partitioned-PDIS), as shown in Figure 7.3. Note this did not require that the primitive policy used options: we merely used the fact that if there are subgoals in the domain where the agent is likely to go through with a fixed state distribution, we can leverage that to decompose the value of a long horizon into the sum over multiple shorter ones. Options is one common way this will occur, but as we see in this example, this can also occur in other ways.

8. Conclusion

We have shown the benefits that come from using options-based behavior and evaluation policies. Options-based policies are more structured allowing lower mean squared error when using importance sampling. Furthermore, special cases may naturally arise due to using options, such as when options terminate in a fixed state distribution; this special case can greatly reduce the mean squared error of importance sampling. Even if you start with a primitive behavior policy, it can still be reinterpreted as an options-based policy and some of the benefits can still apply.

We examine options as a first step, but in the future these results may be extended to full hierarchical policies (like the MAX-Q framework), which are even more powerful. They may also lead to more types of naturally occurring special cases.

Acknowledgments

The research reported here was supported in part by ONR Young Investigator award, an NSF CAREER award and by the Institute of Education Sciences, U.S. Department of Education, through Grants R305A130215 and R305B150008 to Carnegie Mellon University. The opinions expressed are those of the authors and do not represent views of NSF, IES or the U.S. Dept. of Education.

References

Brunskill, Emma and Li, Lihong. Pac-inspired option discovery in lifelong reinforcement learning. In *ICML*, pp. 316–324, 2014.

Dietterich, Thomas G. Hierarchical reinforcement learning with the maxq value function decomposition. *J. Artif. Intell. Res.(JAIR)*, 13:227–303, 2000.

Dudík, M., Langford, J., and Li, L. Doubly robust policy evaluation and learning. In *Proceedings of the Twenty-Eighth International Conference on Machine Learning*, pp. 1097–1104, 2011.

Hallak, Assaf, Schnitzler, François, Mann, Timothy Arthur, and Mannor, Shie. Off-policy model-based learning under unknown factored dynamics. In *ICML*, pp. 711–719, 2015a.

Hallak, Assaf, Tamar, Aviv, Munos, Rémi, and Mannor, Shie. Generalized emphatic temporal difference learning: Bias-variance analysis. *arXiv preprint arXiv:1509.05172*, 2015b.

Jiang, Nan and Li, Lihong. Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 652–661, 2016.

Mankowitz, Daniel J, Mann, Timothy A, and Mannor, Shie. Time regularized interrupting options. In *Internation Conference on Machine Learning*, 2014.

Mann, Timothy A and Mannor, Shie. The advantage of planning with options. *RLDM 2013*, pp. 9, 2013.

Munos, Rémi, Stepleton, Tom, Harutyunyan, Anna, and Bellemare, Marc. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 1046–1054, 2016.

Precup, Doina. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, pp. 80, 2000.

Precup, Doina, Sutton, Richard S, and Dasgupta, Sanjoy. Off-policy temporal-difference learning with function approximation. In *ICML*, pp. 417–424, 2001.

Sutton, Richard S, Precup, Doina, and Singh, Satinder. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.

Thomas, P. S. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, 2016.